arXiv:1403.3427v1 [math.FA] 13 Mar 2014

Explicit Matrices with the Restricted Isometry Property: Breaking the Square-Root Bottleneck

Dustin G. Mixon

Abstract Matrices with the restricted isometry property (RIP) are of particular interest in compressed sensing. To date, the best known RIP matrices are constructed using random processes, while explicit constructions are notorious for performing at the "square-root bottleneck," i.e., they only accept sparsity levels on the order of the square root of the number of measurements. The only known explicit matrix which surpasses this bottleneck was constructed by Bourgain, Dilworth, Ford, Konyagin and Kutzarova in [5]. This chapter provides three contributions to further the groundbreaking work of Bourgain et al.: (i) we develop an intuition for their matrix construction and underlying proof techniques; (ii) we prove a generalized version of their main result; and (iii) we apply this more general result to maximize the extent to which their matrix construction surpasses the square-root bottleneck.

1 Introduction

A matrix Φ is said to satisfy the (K,δ)-*restricted isometry property (RIP)* if

$$(1-\delta)\|x\|^2 \leq \|\Phi x\|^2 \leq (1+\delta)\|x\|^2$$

for every K-sparse vector x. RIP matrices are useful when compressively sensing signals which are sparse in some known orthonormal basis. Indeed, if there is an orthogonal sparsity matrix Ψ such that every signal of interest x has the property that Ψx is K-sparse, then any such x can be stably reconstructed from measurements of the form $y = Ax$ by minimizing $\|\Psi x\|_1$ subject to the measurements, provided $A\Psi^{-1}$ satisfies $(2K,\delta)$-RIP with $\delta < \sqrt{2}-1$ [8]. For sensing regimes in which measurements are costly, it is desirable to minimize the number of measurements

Dustin G. Mixon
Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio, USA, e-mail: dustin.mixon@afit.edu

necessary for signal reconstruction; this corresponds to the number of rows M in the $M \times N$ sensing matrix A. One can apply the theory of Gelfand widths to show that stable reconstruction by L1-minimization requires $K = O(M/\log(N/M))$ [4], and random matrices show that this bound is essentially tight; indeed, $M \times N$ matrices with iid subgaussian entries satisfy $(2K, \delta)$-RIP with high probability provided $M = \Omega_\delta(K \log(N/K))$ [12].

Unfortunately, random matrices are not always RIP, though the failure rate vanishes asymptotically. In applications, you might wish to verify that your randomly drawn matrix actually satisfies RIP before designing your sensing platform around that matrix, but unfortunately, this is NP-hard in general [2]. As such, one is forced to blindly assume that the randomly drawn matrix is RIP, and admittedly, this is a reasonable assumption considering the failure rate. Still, this is dissatisfying from a theoretical perspective, and it motivates the construction of explicit RIP matrices:

Definition 1. Let ExRIP$[z]$ denote the following statement:

There exists an explicit family of $M \times N$ matrices with arbitrarily large aspect ratio N/M which are (K, δ)-RIP with $K = \Omega(M^{z-\varepsilon})$ for all $\varepsilon > 0$ and $\delta < \sqrt{2} - 1$.

Since there exist (non-explicit) matrices satisfying $z = 1$ above, the goal is to prove ExRIP$[1]$. The most common way to demonstrate that an explicit matrix Φ satisfies RIP is to leverage the pairwise incoherence between the columns of Φ. Indeed, it is straightforward to prove ExRIP$[1/2]$ by taking Φ to have near-optimally incoherent unit-norm columns and appealing to interpolation of operators or Gershgorin's circle theorem (e.g., see [1, 10, 11]). The emergence of this "square-root bottleneck" compelled Tao to pose the explicit construction of RIP matrices as an open problem [18]. Since then, only one construction has managed to break the bottleneck: In [5], Bourgain, Dilworth, Ford, Konyagin and Kutzarova prove ExRIP$[1/2 + \varepsilon_0]$ for some undisclosed $\varepsilon_0 > 0$. This constant has since been estimated as $\varepsilon_0 \approx 5.5169 \times 10^{-28}$ [14].

Instead of estimating δ in terms of coherence, Bourgain et al. leverage additive combinatorics to construct Φ and to demonstrate certain cancellations in the Gram matrix $\Phi^*\Phi$. Today (three years later), this is the only known explicit construction which breaks the square-root bottleneck, thereby leading to two natural questions:

- What are the proof techniques that Bourgain et al. applied?
- Can we optimize the analysis to increase ε_0?

These questions were investigated recently in a series of blog posts [14, 15, 16], on which this chapter is based. In the next section, we provide some preliminaries—we first cover the techniques used in [5] to demonstrate RIP, and then we discuss some basic additive combinatorics to motivate the matrix construction. Section 3 then describes the construction of Φ, namely a subcollection of the chirps studied in [9], and discusses one method of selecting chirps (i.e., the method of Bourgain et al.). Section 4 provides the main result, namely *the BDFKK restricted isometry machine*, which says that a "good" selection of chirps will result in an RIP matrix construction which breaks the square-root bottleneck. This is a generalization of the main result in [5], as it offers more general sufficient conditions for good chirp selection, but the

proof is similar. After generalizing the sufficient conditions, we optimize over these conditions to increase the largest known ε_0 for which ExRIP$[1/2+\varepsilon_0]$ holds:

$$\varepsilon_0 \approx 4.4466 \times 10^{-24}.$$

Of course, any improvement to the chirp selection method will further increase this constant, and hopefully, the BDFKK restricted isometry machine and overall intuition provided in this chapter will foster such progress. Section 5 contains the proofs of certain technical lemmas that are used to prove the main result.

2 Preliminaries

The goal of this section is to provide some intuition for the main ideas in [5]. We first explain the overall proof technique for demonstrating RIP (this is the vehicle for breaking the square-root bottleneck), and then we introduce some basic ideas from additive combinatorics.

2.1 The Big-Picture Techniques

Before explaining how Bourgain et al. broke the square-root bottleneck, let's briefly discuss the more common, coherence-based technique to demonstrate RIP. Let $\Phi_{\mathscr{K}}$ denote the submatrix of Φ whose columns are indexed by $\mathscr{K} \subseteq \{1,\dots,N\}$. Then (K,δ)-RIP equivalently states that, for every $\mathscr{K}$ of size K, the eigenvalues of $\Phi_{\mathscr{K}}^{*}\Phi_{\mathscr{K}}$ lie in $[1-\delta,1+\delta]$. As such, we can prove that a matrix is RIP by approximating eigenvalues. To this end, if we assume the columns of Φ have unit norm, and if we let μ denote the largest off-diagonal entry of $\Phi^{*}\Phi$ in absolute value (this is the worst-case coherence of the columns of Φ), then the Gershgorin circle theorem implies that Φ is $(K,(K-1)\mu)$-RIP. Unfortunately, the coherence can't be too small, due to the Welch bound [17]:

$$\mu \geq \sqrt{\frac{N-M}{M(N-1)}},$$

which is $\Omega(M^{-1/2})$ provided $N \geq cM$ for some $c>1$. Thus, to get $(K-1)\mu = \delta < 1/2$, we require $K < 1/(2\mu)+1 = O(M^{1/2})$. This is much smaller than the random RIP constructions which instead take $K = O(M^{1-\varepsilon})$ for all $\varepsilon > 0$, thereby revealing the shortcoming of the Gershgorin technique.

Now let's discuss the alternative techniques that Bourgain et al. use. The main idea is to convert the RIP statement, which concerns all K-sparse vectors simultaneously, into a statement about finitely many vectors:

Definition 2 (flat RIP). We say $\Phi = [\varphi_1 \cdots \varphi_N]$ satisfies (K,θ)-*flat RIP* if for every disjoint $I,J \subseteq \{1,\ldots,N\}$ of size $\leq K$,

$$\left|\left\langle \sum_{i\in I} \varphi_i, \sum_{j\in J} \varphi_j \right\rangle\right| \leq \theta\sqrt{|I||J|}.$$

Lemma 1 (essentially Lemma 3 in [5], cf. Theorem 13 in [3]). *If Φ has (K,θ)-flat RIP and unit-norm columns, then Φ has $(K, 150\theta\log K)$-RIP.*

Unlike the coherence argument, flat RIP doesn't lead to much loss in K. In particular, [3] shows that random matrices satisfy (K,θ)-flat RIP with $\theta = O(\delta/\log K)$ when $M = \Omega((K/\delta^2)\log^2 K\log N)$. As such, it makes sense that flat RIP would be a vehicle to break the square-root bottleneck. However, in practice, it's difficult to control both the left- and right-hand sides of the flat RIP inequality—it would be much easier if we only had to worry about getting cancellations, and not getting different levels of cancellation for different-sized subsets. This leads to the following:

Definition 3 (weak flat RIP). We say $\Phi = [\varphi_1 \cdots \varphi_N]$ satisfies (K,θ')-*weak flat RIP* if for every disjoint $I,J \subseteq \{1,\ldots,N\}$ of size $\leq K$,

$$\left|\left\langle \sum_{i\in I} \varphi_i, \sum_{j\in J} \varphi_j \right\rangle\right| \leq \theta' K.$$

Lemma 2 (essentially Lemma 1 in [5]). *If Φ has (K,θ')-weak flat RIP and worst-case coherence $\mu \leq 1/K$, then Φ has $(K,\sqrt{\theta'})$-flat RIP.*

Proof. By the triangle inequality, we have

$$\left|\left\langle \sum_{i\in I} \varphi_i, \sum_{j\in J} \varphi_j \right\rangle\right| \leq \sum_{i\in I}\sum_{j\in J} |\langle \varphi_i, \varphi_j\rangle| \leq |I||J|\mu \leq |I||J|/K.$$

Since Φ also has weak flat RIP, we then have

$$\left|\left\langle \sum_{i\in I} \varphi_i, \sum_{j\in J} \varphi_j \right\rangle\right| \leq \min\{\theta' K, |I||J|/K\} \leq \sqrt{\theta'|I||J|}. \quad \square$$

Unfortunately, this coherence requirement puts K back in the square-root bottleneck, since $\mu \leq 1/K$ is equivalent to $K \leq 1/\mu = O(M^{1/2})$. To rectify this, Bourgain et al. use a trick in which a modest K with tiny δ can be converted to a large K with modest δ:

Lemma 3 (buried in Lemma 3 in [5], cf. Theorem 1 in [13]). *If Φ has (K,δ)-RIP, then Φ has $(sK, 2s\delta)$-RIP for all $s \geq 1$.*

In [13], this trick is used to get RIP results for larger K when testing RIP for smaller K. For the explicit RIP matrix problem, we are stuck with proving how small δ is when K on the order of $M^{1/2}$. Note that this trick will inherently exhibit

some loss in K. Assuming the best possible scaling for all N, K and δ is $M = \Theta((K/\delta^2)\log(N/K))$, then if $N = \text{poly}(M)$, you can get $(M^{1/2}, \delta)$-RIP only if $\delta = \Omega((\log^{1/2} M)/M^{1/4})$. In this best-case scenario, you would want to pick $s = M^{1/4-\varepsilon}$ for some $\varepsilon > 0$ and apply Lemma 3 to get $K = O(M^{3/4-\varepsilon})$. In some sense, this is another manifestation of the square-root bottleneck, but it would still be a huge achievement to saturate this bound.

2.2 A Brief Introduction to Additive Combinatorics

In this subsection, we briefly detail some key ideas from additive combinatorics; the reader is encouraged to see [19] for a more complete introduction. Given an additive group G and finite sets $A, B \subseteq G$, we can define the sumset

$$A+B := \{a+b : a \in A,\ b \in B\},$$

the difference set

$$A-B := \{a-b : a \in A,\ b \in B\},$$

and the additive energy

$$E(A,B) := \#\big\{(a_1,a_2,b_1,b_2) \in A^2 \times B^2 : a_1+b_1 = a_2+b_2\big\}.$$

These definitions are useful in quantifying the additive structure of a set. In particular, consider the following:

Lemma 4. *A nonempty subset A of some additive group G satisfies the following inequalities:*

(i) $|A+A| \geq |A|$
(ii) $|A-A| \geq |A|$
(iii) $E(A,A) \leq |A|^3$

with equality precisely when A is a translate of some subgroup of G.

Proof. For (i), pick $a \in A$. Then $|A+A| \geq |A+a| = |A|$. Considering

$$A+A = \bigcup_{a \in A} (A+a),$$

we have equality in (i) precisely when $A+A = A+a$ for every $a \in A$. Equivalently, given $a_0 \in A$, then for every $a \in A$, addition by $a-a_0$ permutes the members of $A+a_0$. This is further equivalent to the following: Given $a_0 \in A$, then for every $a \in A$, addition by $a-a_0$ permutes the members of $A-a_0$. It certainly suffices for $H := A-a_0$ to be a group, and it is a simple exercise to verify that this is also necessary.

The proof for (ii) is similar.

For (iii), we note that

$$E(A,A) = \#\big\{(a,b,c) \in A^3 : a+b-c \in A\big\} \leq |A|^3,$$

with equality precisely when A has the property that $a+b-c \in A$ for every $a,b,c \in A$. Again, it clearly suffices for $A - a_0$ to be a group, and necessity is a simple exercise. $\square$

The notion of additive structure is somewhat intuitive. You should think of a translate of a subgroup as having maximal additive structure. When the bounds (i), (ii) and (iii) are close to being achieved by A (e.g., A is an arithmetic progression), you should think of A as having a lot of additive structure. Interestingly, while there are different measures of additive structure (e.g., $|A-A|$ and $E(A,A)$), they often exhibit certain relationships (perhaps not surprisingly). The following is an example of such a relationship which is used throughout the paper by Bourgain et al. [5]:

Lemma 5 (Corollary 1 in [5]). *If $E(A,A) \geq |A|^3/K$, then there exists a set $A' \subseteq A$ such that $|A'| \geq |A|/(20K)$ and $|A'-A'| \leq 10^7 K^9 |A|$.*

In words, a set with a lot of additive energy necessarily has a large subset with a small difference set. This is proved using a version of the Balog–Szemeredi–Gowers lemma [6].

If translates of subgroups have maximal additive structure, then which sets have minimal additive structure? It turns out that random sets tend to (nearly) have this property, and one way to detect low additive structure is Fourier bias:

$$\|A\|_u := \max_{\substack{\theta \in G \\ \theta \neq 0}} |\widehat{1_A}(\theta)|,$$

where the Fourier transform ($\widehat{\cdot}$) used here has a $1/|G|$ factor in front (it is not unitary). For example, if $G = \mathbb{Z}/n\mathbb{Z}$, we take

$$\hat{f}(\xi) := \frac{1}{|G|} \sum_{x \in G} f(x) e^{-2\pi i x\xi/n}.$$

Interestingly, $\|A\|_u$ captures how far $E(A,A)$ is from its minimal value $|A|^4/|G|$:

Lemma 6. *For any subset A of a finite additive group G, we have*

(i) $E(A,A) \geq \frac{|A|^4}{|G|}$, *and*

(ii) $\|A\|_u^4 \leq \frac{1}{|G|^3}\Big(E(A,A) - \frac{|A|^4}{|G|}\Big) \leq \frac{|A|}{|G|}\|A\|_u^2$.

Proof. Define $\lambda_x := \#\{(a,a') \in A^2 : a - a' = x\}$. Then (i) follows from Cauchy–Schwarz:

$$|A|^2 = \sum_{x \in G} \lambda_x \leq |G|^{1/2} \|\lambda\|_2 = \big(|G|E(A,A)\big)^{1/2}.$$

We will prove (ii) assuming $G = \mathbb{Z}/n\mathbb{Z}$, but the proof generalizes. Denote $e_n(x) := e^{2\pi i x/n}$. For the left-hand inequality, we consider

$$\sum_{\theta\in G}|\widehat{1_A}(\theta)|^4=\sum_{\theta\in G}\left|\frac{1}{|G|}\sum_{a\in A}e_n(-\theta a)\right|^4=\frac{1}{|G|^4}\sum_{\theta\in G}\left|\sum_{x\in G}\lambda_x e_n(-\theta x)\right|^2,$$

where the last step is by expanding $|w|^2=w\overline{w}$. Then Parseval's identity simplifies this to $\frac{1}{|G|^3}\|\lambda\|_2^2=\frac{1}{|G|^3}E(A,A)$. We use this to bound $\|A\|_u^4$:

$$\|A\|_u^4=\max_{\substack{\theta\in G\\ \theta\neq 0}}|\widehat{1_A}(\theta)|^4\leq\sum_{\substack{\theta\in G\\ \theta\neq 0}}|\widehat{1_A}(\theta)|^4=\frac{1}{|G|^3}E(A,A)-\frac{|A|^4}{|G|^4}.$$

For the right-hand inequality, we apply Parseval's identity:

$$E(A,A)=\sum_{x\in G}\lambda_x^2=\frac{1}{|G|}\sum_{\theta\in G}\left|\sum_{x\in G}\lambda_x e_n(-\theta x)\right|^2=\frac{|A|^4}{|G|}+\frac{1}{|G|}\sum_{\substack{\theta\in G\\ \theta\neq 0}}\left|\sum_{x\in G}\lambda_x e_n(-\theta x)\right|^2$$

From here, we apply the expansion $|w|^2=w\overline{w}$

$$\left|\sum_{a\in A}e_n(-\theta a)\right|^2=\sum_{x\in G}\lambda_x e_n(-\theta x)$$

to continue:

$$\sum_{\substack{\theta\in G\\ \theta\neq 0}}\left|\sum_{x\in G}\lambda_x e_n(-\theta x)\right|^2=\sum_{\substack{\theta\in G\\ \theta\neq 0}}\left|\sum_{a\in A}e_n(-\theta a)\right|^4\leq\sum_{\substack{\theta\in G\\ \theta\neq 0}}\big(|G|\|A\|_u\big)^2\left|\sum_{a\in A}e_n(-\theta a)\right|^2.$$

Applying Parseval's identity then gives

$$E(A,A)\leq\frac{|A|^4}{|G|}+\big(|G|\|A\|_u\big)^2\cdot\frac{1}{|G|}\sum_{\theta\in G}\left|\sum_{a\in A}e_n(-\theta a)\right|^2=\frac{|A|^4}{|G|}+|G|^2\|A\|_u^2|A|,$$

which is a rearrangement of the right-hand inequality. □

3 The Matrix Construction

This section combines ideas from the previous section to introduce the matrix construction used by Bourgain et al. [5] to break the square-root bottleneck. The main idea is to construct a Gram matrix $\Phi^*\Phi$ whose entries exhibit cancellations for weak flat RIP (see Definition 3). By Lemma 6, we can control cancellations of complex exponentials

$$\left|\sum_{a\in A}e_n(\theta a)\right|\leq n\|A\|_u,\qquad\theta\neq 0$$

in terms of the additive energy of the index set $A \subseteq \mathbb{Z}/n\mathbb{Z}$; recall that $e_n(x) := e^{2\pi ix/n}$. This motivates us to pursue a Gram matrix whose entries are complex exponentials. To this end, consider the following vector:

$$u_{a,b} := \frac{1}{\sqrt{p}}\Big(e_p(ax^2+bx)\Big)_{x\in\mathbb{F}_p},$$

where p is prime and $\mathbb{F}_p$ denotes the field of size p. Such vectors are called *chirps*, and they are used in a variety of applications including radar. Here, we are mostly interested in the form of their inner products. If $a_1 = a_2$, then $\langle u_{a_1,b_1}, u_{a_2,b_2}\rangle = \delta_{b_1,b_2}$ by the geometric sum formula. Otherwise, the inner product is more interesting:

$$\langle u_{a_1,b_1}, u_{a_2,b_2}\rangle = \frac{1}{p}\sum_{x\in\mathbb{F}_p} e_p\Big((a_1-a_2)x^2+(b_1-b_2)x\Big).$$

Since $a_1-a_2\neq 0$, we can complete the square in the exponent, and changing variables to $y := x+(b_1-b_2)/(2(a_1-a_2))$ gives

$$\langle u_{a_1,b_1}, u_{a_2,b_2}\rangle = \frac{1}{p}e_p\bigg(-\frac{(b_1-b_2)^2}{4(a_1-a_2)}\bigg)\sum_{y\in\mathbb{F}_p} e_p\Big((a_1-a_2)y^2\Big).$$

Finally, this can be simplified using a quadratic Gauss sum formula:

$$\langle u_{a_1,b_1}, u_{a_2,b_2}\rangle = \frac{\sigma_p}{\sqrt{p}}\bigg(\frac{a_1-a_2}{p}\bigg)e_p\bigg(-\frac{(b_1-b_2)^2}{4(a_1-a_2)}\bigg),$$

where σ_p is 1 or i (depending on whether p is 1 or 3 mod 4) and $(\frac{a_1-a_2}{p})$ is a Legendre symbol, taking value ± 1 depending on whether a_1-a_2 is a perfect square mod p. Modulo these factors, the above inner product is a complex exponential, and since we want these in our Gram matrix $\Phi^*\Phi$, we will take Φ to have columns of the form $u_{a,b}$—in fact, the columns will be $\{u_{a,b}\}_{(a,b)\in\mathscr{A}\times\mathscr{B}}$ for some well-designed sets $\mathscr{A},\mathscr{B}\subseteq\mathbb{F}_p$.

For weak flat RIP, we want to bound the following quantity for every $\Omega_1,\Omega_2\subseteq\mathscr{A}\times\mathscr{B}$ with $|\Omega_1|,|\Omega_2|\leq\sqrt{p}$:

$$\bigg|\bigg\langle \sum_{(a_1,b_1)\in\Omega_1} u_{a_1,b_1}, \sum_{(a_2,b_2)\in\Omega_2} u_{a_2,b_2}\bigg\rangle\bigg|.$$

For $i=1,2$, define

$$A_i := \{a : \exists b \text{ s.t. } (a,b)\in\Omega_i\} \quad\text{and}\quad \Omega_i(a) := \{b : (a,b)\in\Omega_i\}.$$

These provide an alternate expression for the quantity of interest:

$$\begin{aligned}\Big|\sum_{(a_1,b_1)\in\Omega_1}\sum_{(a_2,b_2)\in\Omega_2}\langle u_{a_1,b_1},u_{a_2,b_2}\rangle\Big| &= \Big|\sum_{\substack{a_1\in A_1\\ a_2\in A_2}}\sum_{\substack{b_1\in\Omega_1(a_1)\\ b_2\in\Omega_2(a_2)}}\langle u_{a_1,b_1},u_{a_2,b_2}\rangle\Big|,\\ &\le \sum_{\substack{a_1\in A_1\\ a_2\in A_2}}\Big|\sum_{\substack{b_1\in\Omega_1(a_1)\\ b_2\in\Omega_2(a_2)}}\langle u_{a_1,b_1},u_{a_2,b_2}\rangle\Big|\\ &= \frac{1}{\sqrt{p}}\sum_{\substack{a_1\in A_1\\ a_2\in A_2}}\Big|\sum_{\substack{b_1\in\Omega_1(a_1)\\ b_2\in\Omega_2(a_2)}}e_p\Big(-\frac{(b_1-b_2)^2}{4(a_1-a_2)}\Big)\Big|.\end{aligned}$$

Pleasingly, it now suffices to bound a sum of complex exponentials, which we feel equipped to do using additive combinatorics. The following lemma does precisely this (it can be viewed as an analog of Lemma 6).

Lemma 7 (Lemma 9 in [5]). *For every $\theta\in\mathbb{F}_p^*$ and $B_1,B_2\subseteq\mathbb{F}_p$, we have*

$$\Big|\sum_{\substack{b_1\in B_1\\ b_2\in B_2}}e_p\Big(\theta(b_1-b_2)^2\Big)\Big|\le |B_1|^{1/2}E(B_1,B_1)^{1/8}|B_2|^{1/2}E(B_2,B_2)^{1/8}p^{1/8}.$$

Proof. First, Cauchy–Schwarz gives

$$\begin{aligned}\Big|\sum_{\substack{b_1\in B_1\\ b_2\in B_2}}e_p\Big(\theta(b_1-b_2)^2\Big)\Big|^2 &= \Big|\sum_{b_1\in B_1}1\cdot\sum_{b_2\in B_2}e_p\Big(\theta(b_1-b_2)^2\Big)\Big|^2\\ &\le |B_1|\sum_{b_1\in B_1}\Big|\sum_{b_2\in B_2}e_p\Big(\theta(b_1-b_2)^2\Big)\Big|^2.\end{aligned}$$

Expanding $|w|^2=w\overline{w}$ and rearranging then gives an alternate expression for the right-hand side:

$$|B_1|\sum_{b_2,b_2'\in B_2}e_p\Big(\theta(b_2^2-(b_2')^2)\Big)\overline{\sum_{b_1\in B_1}e_p\Big(\theta(2b_1(b_2-b_2'))\Big)}.$$

Applying Cauchy–Schwarz again, we then have

$$\Big|\sum_{\substack{b_1\in B_1\\ b_2\in B_2}}e_p\Big(\theta(b_1-b_2)^2\Big)\Big|^4\le |B_1|^2|B_2|^2\sum_{b_2,b_2'\in B_2}\Big|\sum_{b_1\in B_1}e_p\Big(\theta(2b_1(b_2-b_2'))\Big)\Big|^2,$$

and expanding $|w|^2=w\overline{w}$ this time gives

$$|B_1|^2|B_2|^2\sum_{\substack{b_1,b_1'\in B_1\\ b_2,b_2'\in B_2}}e_p\Big(2\theta(b_1-b_1')(b_2-b_2')\Big).$$

At this point, it is convenient to change variables, namely, $x = b_1 - b_1'$ and $y = b_2 - b_2'$:

$$\left| \sum_{\substack{b_1 \in B_1 \\ b_2 \in B_2}} e_p\Big(\theta(b_1 - b_2)^2\Big) \right|^4 \leq |B_1|^2 |B_2|^2 \sum_{x,y \in \mathbb{F}_p} \lambda_x \mu_y e_p(2\theta xy), \tag{1}$$

where $\lambda_x := \#\{(b_1, b_1') \in B_1^2 : b_1 - b_1' = x\}$ and similarly for μ_y in terms of B_2. We now apply Cauchy–Schwarz again to bound the sum in (1):

$$\left| \sum_{x \in \mathbb{F}_p} \lambda_x \sum_{y \in \mathbb{F}_p} \mu_y e_p(2\theta xy) \right|^2 \leq \|\lambda\|_2^2 \sum_{x \in \mathbb{F}_p} \left| \sum_{y \in \mathbb{F}_p} \mu_y e_p(2\theta xy) \right|^2,$$

and changing variables $x' := -2\theta x$ (this change is invertible since $\theta \neq 0$), we see that the right-hand side is a sum of squares of the Fourier coefficients of μ. As such, Parseval's identity gives the following simplification:

$$\left| \sum_{x,y \in \mathbb{F}_p} \lambda_x \mu_y e_p(2\theta xy) \right|^2 \leq p \|\lambda\|_2^2 \|\mu\|_2^2 = pE(B_1, B_1)E(B_2, B_2).$$

Applying this bound to (1) gives the result. □

3.1 How to Construct $\mathscr{B}$

Lemma 7 enables us to prove weak-flat-RIP-type cancellations in cases where $\Omega_1(a_1), \Omega_2(a_2) \subseteq \mathscr{B}$ both lack additive structure. Indeed, the method of [5] is to do precisely this, and the remaining cases (where either $\Omega_1(a_1)$ or $\Omega_2(a_2)$ has more additive structure) will find cancellations by accounting for the dilation weights $1/(a_1 - a_2)$. Overall, we will be very close to proving that Φ is RIP if most subsets of $\mathscr{B}$ lack additive structure. To this end, Bourgain et al. [5] actually prove something much stronger: They design $\mathscr{B}$ in such a way that all sufficiently large subsets have low additive structure. The following theorem is the first step in the design:

Theorem 1 (Theorem 5 in [5]). *Fix $r, M \in \mathbb{N}$, $M \geq 2$, and define the cube $\mathscr{C} := \{0, \ldots, M-1\}^r \subseteq \mathbb{Z}^r$. Let τ denote the solution to the equation*

$$\Big(\frac{1}{M}\Big)^{2\tau} + \Big(\frac{M-1}{M}\Big)^{\tau} = 1.$$

Then for any subsets $A, B \subseteq \mathscr{C}$, we have

$$|A + B| \geq \big(|A||B|\big)^{\tau}.$$

As a consequence of this theorem (taking $A = B$), we have $|A+A| \geq |A|^{2\tau}$ for every $A \subseteq \mathscr{C}$, and since $\tau > 1/2$, this means that large subsets A have $|A+A| \gg |A|$, indicating low additive structure. However, $\mathscr{C}$ is a subset of the group $\mathbb{Z}^r$, whereas we need to construct a subset $\mathscr{B}$ of $\mathbb{F}_p$. The trick here is to pick $\mathscr{B}$ so that it inherits the additive structure of $\mathscr{C}$, and we use a *Freiman isomorphism* to accomplish this. In particular, we want a mapping $\varphi\colon \mathscr{C} \to \mathbb{F}_p$ such that $c_1 + c_2 = c_3 + c_4$ if and only if $\varphi(c_1) + \varphi(c_2) = \varphi(c_3) + \varphi(c_4)$, and we will take $\mathscr{B} := \varphi(\mathscr{C})$—this is what it means for $\mathscr{C}$ and $\mathscr{B}$ to be Freiman isomorphic, and it's easy to see that Freiman isomorphic sets have the same sized sumsets, difference sets and additive energy. In this case, it suffices to take

$$\mathscr{B} := \left\{ \sum_{j=1}^{r} x_j (2M)^{j-1} : x_1, \ldots, x_r \in \{0, \ldots, M-1\} \right\}. \tag{2}$$

Indeed, the $2M$-ary expansion of $b_1, b_2 \in \mathscr{B}$ reveals the $c_1, c_2 \in \mathscr{C}$ such that $\varphi(c_1) = b_1$ and $\varphi(c_2) = b_2$. Also, adding b_1 and b_2 incurs no carries, so the expansion of $b_1 + b_2$ reveals $c_1 + c_2$ (even when $c_1 + c_2 \notin \mathscr{C}$).

We already know that large subsets of $\mathscr{C}$ (and $\mathscr{B}$) exhibit low additive structure, but the above theorem only gives this in terms of the sumset, whereas Lemma 7 requires low additive structure in terms of additive energy. As such, we will first convert the above theorem into a statement about difference sets, and then apply Lemma 5 to further convert it in terms of additive energy:

Corollary 1 (essentially Corollary 3 in [5]). *Fix r, M and τ according to Theorem 1, take $\mathscr{B}$ as defined in* (2)*, and pick s and t such that $(2\tau - 1)s \geq t$. Then every subset $B \subseteq \mathscr{B}$ such that $|B| > p^s$ satisfies $|B - B| > p^t |B|$.*

Proof. First note that $-B$ is a translate of some other set $B' \subseteq \mathscr{B}$. Explicitly, if $b_0 = \sum_{j=1}^{r} (M-1)(2M)^{j-1}$, then we can take $B' := b_0 - B$. As such, Theorem 1 gives

$$|B - B| = |B + B'| \geq |B|^{2\tau} = |B|^{2\tau - 1} |B| > p^{(2\tau-1)s} |B| \geq p^t |B|. \quad \square$$

Corollary 2 (essentially Corollary 4 in [5]). *Fix r, M and τ according to Theorem 1, take $\mathscr{B}$ as defined in* (2)*, and pick γ and ℓ such that $(2\tau - 1)(\ell - \gamma) \geq 10\gamma$. Then for every $\varepsilon > 0$, there exists P such that for every $p \geq P$, every subset $S \subseteq \mathscr{B}$ with $|S| > p^\ell$ satisfies $E(S,S) < p^{-\gamma+\varepsilon} |S|^3$.*

Proof. Suppose to the contrary that there exists $\varepsilon > 0$ such that there are arbitrarily large p for which there is a subset $S \subseteq \mathscr{B}$ with $|S| > p^\ell$ and $E(S,S) \geq p^{-\gamma+\varepsilon} |S|^3$. Writing $E(S,S) = |S|^3/K$, then $K \leq p^{\gamma-\varepsilon}$. By Lemma 5, there exists $B \subseteq S$ such that, for sufficiently large p,

$$|B| \geq |S|/(20K) > \frac{1}{20} p^{\ell-\gamma+\varepsilon} > p^{\ell-\gamma},$$

and

$$|B-B| \leq 10^7 K^9 |S| \leq 10^7 K^9 (20K|B|) \leq 10^7 \cdot 20 \cdot p^{10(\gamma-\varepsilon)}|B| < p^{10\gamma}|B|.$$

However, this contradicts the previous corollary with $s = \ell - \gamma$ and $t = 10\gamma$. □

Notice that we can weaken our requirements on γ and ℓ if we had a version of Lemma 5 with a smaller exponent on K. This exponent comes from a version of the Balog–Szemeredi–Gowers lemma (Lemma 6 in [5]), which follows from the proof of Lemma 2.2 in [6]. (Specifically, take $A = B$, and you need to change $A -_E B$ to $A +_E B$, but this change doesn't affect the proof.) Bourgain et al. indicate that it would be desirable to prove a better version of this lemma, but it is unclear how easy that would be.

3.2 How to Construct $\mathscr{A}$

The previous subsection showed how to construct $\mathscr{B}$ so as to ensure that all sufficiently large subsets have low additive structure. By Lemma 7, this in turn ensures that Φ exhibits weak-flat-RIP-type cancellations for most $\Omega_1(a_1), \Omega_2(a_2) \subseteq \mathscr{B}$. For the remaining cases, Φ must exhibit weak-flat-RIP-type cancellations by somehow leveraging properties of $\mathscr{A}$.

The next section gives the main result, which requires a subset $\mathscr{A} = \mathscr{A}(p)$ of $\mathbb{F}_p$ for which there exists an even number m as well as an $\alpha > 0$ (both independent of p) such that the following two properties hold:

(i) $\Omega(p^\alpha) \leq |\mathscr{A}(p)| \leq p^\alpha$.

(ii) For each $a \in \mathscr{A}$, then $a_1, \ldots, a_{2m} \in \mathscr{A} \setminus \{a\}$ satisfies

$$\sum_{j=1}^{m} \frac{1}{a-a_j} = \sum_{j=m+1}^{2m} \frac{1}{a-a_j} \tag{3}$$

only if $(a_1, \ldots, a_m)$ and $(a_{m+1}, \ldots, a_{2m})$ are permutations of each other. Here, division (and addition) is taken in the field $\mathbb{F}_p$.

Unfortunately, these requirements on $\mathscr{A}$ lead to very little intuition compared to our current understanding of $\mathscr{B}$. Regardless, we will continue by considering how Bourgain et al. constructs $\mathscr{A}$. The following lemma describes their construction and makes a slight improvement to the value of α chosen in [5]:

Lemma 8. *Pick $\varepsilon > 0$ and take $L := \lfloor p^{1/2m(4m-1)} \rfloor$ and $U := \lfloor L^{4m-1} \rfloor$. Then*

$$\mathscr{A} := \{x^2 + Ux : 1 \leq x \leq L\}$$

satisfies (i) and (ii) above if we take

$$\alpha = \frac{1}{2m(4m-1)}$$

and p is a sufficiently large prime.

Proof. One may quickly verify (i). For (ii), we claim it suffices to show that for any $n \in \{1,\dots,2m\}$, any distinct $x, x_1, \dots, x_n \in \{1,\dots,L\}$, and any nonzero integers $\lambda_1, \dots, \lambda_n$ such that $|\lambda_1| + \cdots + |\lambda_n| \leq 2m$, then

$$V = \sum_{j=1}^{n} \frac{\lambda_j}{(x-x_j)(x+x_j+U)} \tag{4}$$

is nonzero (in $\mathbb{F}_p$). To see this, define $a := x^2 + Ux$ and $a_j := x_j^2 + Ux_j$. Then

$$V = \sum_{j=1}^{n} \frac{\lambda_j}{a-a_j}$$

As such, if (ii) fails to hold, then subtracting the right-hand side of (3) from the left-hand side produces an example of V which is zero, violating our statement. Thus, our statement implies (ii) by the contrapositive.

We now seek to prove our statement. To this end, define $D_1 := \prod_{j=1}^{n}(x-x_j)$ and $D_2 := \prod_{j=1}^{n}(x+x_j+U)$. Note that (4) being nonzero in $\mathbb{F}_p$ is equivalent to having p not divide $D_1 D_2 V$ as an integer. To prove this, we will show that

(a) p does not divide D_1,
(b) $D_2 V$ is nonzero, and
(c) $|D_2 V| < p$.

Indeed, (b) and (c) together imply that p does not divide $D_2 V$, which combined with (a) implies that p does not divide $D_1 D_2 V$.

For (a), we have $D_1 \neq 0$ since x and the x_j's are distinct by assumption, and since x and each x_j has size at most L, we also have $|D_1| \leq L^{2m}$. To complete the proof of (a), it then suffices to have

$$L^{2m} < p, \tag{5}$$

which we will verify later.

For (b), we will prove that $D_1 D_2 V$ is nonzero. We first write

$$D_1 D_2 V = \sum_{j=1}^{n} \frac{\lambda_j D_1}{x-x_j} \frac{D_2}{x+x_j+U}.$$

For each term in the above sum, note that both fractions are integers, and for every $j \neq 1$, $x+x_1+U$ is a factor of $D_2/(x+x_j+U)$. As such, the entire sum is congruent to the first term modulo $x+x_1+U$:

$$D_1 D_2 V \equiv \lambda_1 \prod_{j=2}^{n}(x-x_j) \prod_{j=2}^{n}(x+x_j+U) \mod (x+x_1+U).$$

Each factor of the form $x+x_j+U$ can be further simplified:

$$x+x_j+U = (x_j-x_1)+(x+x_1+U) \equiv x_j - x_1 \mod (x+x_1+U),$$

and so

$$D_1 D_2 V \equiv V_1 \mod (x + x_1 + U),$$

where

$$V_1 := \lambda_1 \prod_{j=2}^{n} (x - x_j) \prod_{j=2}^{n} (x_j - x_1).$$

To prove that $D_1 D_2 V$ is nonzero, it suffices to show that $x + x_1 + U$ does not divide V_1. To this end, we first note that V_1 is nonzero since x and the x_j's are distinct by assumption. Next, since $|\lambda_1| \leq 2m$ and x and each x_j is at most L, we have $|V_1| \leq 2mL^{2n-2} \leq 2mL^{4m-2}$, and so it suffices to have

$$2mL^{4m-2} \leq U \tag{6}$$

since $U < x + x_1 + U$. We will verify this later.

For (c), we have

$$|D_2 V| \leq \sum_{j=1}^{n} \frac{|\lambda_j|}{|x - x_j|} \prod_{\substack{j'=1 \\ j' \neq j}}^{n} |x + x_{j'} + U| \leq \left(\sum_{j=1}^{n} |\lambda_j| \right) \cdot (2L + U)^{n-1}.$$

Considering our assumption on the λ_j's we then have $|D_2 V| \leq 2m(2L + U)^{n-1} \leq 2m(2L + U)^{2m-1}$, and so it suffices to have

$$2m(2L + U)^{2m-1} < p. \tag{7}$$

Overall, we have shown it suffices to have (5), (6) and (7). To satisfy (6) for sufficiently large L, we take $U := \lfloor L^{4m-2+\varepsilon} \rfloor$. Then $L = o(U)$, and so $2m(2L + U)^{2m-1} < U^{2m-1+\varepsilon} \leq L^{(4m-2+\varepsilon)(2m-1+\varepsilon)}$ for sufficiently large U. As such, for (7), it suffices to take $L := \lfloor p^{1/(4m-2+\varepsilon)(2m-1+\varepsilon)} \rfloor$, which also satisfies (5). For simplicity, we take $\varepsilon = 1$. $\square$

4 The Main Result

We are now ready to state the main result of this chapter, which is a generalization of the main result in [5]. Later in this section, we will maximize ε_0 such that this result implies ExRIP$[1/2 + \varepsilon_0]$ with the matrix construction from [5].

Theorem 2 (The BDFKK restricted isometry machine). *For every prime p, define subsets $\mathscr{A} = \mathscr{A}(p)$ and $\mathscr{B} = \mathscr{B}(p)$ of $\mathbb{F}_p$. Suppose there exist constants $m \in 2\mathbb{N}$, $\ell, \gamma > 0$ independent of p such that the following conditions apply:*

(a) For every sufficiently large p, and for every $a \in \mathscr{A}$ and $a_1, \ldots, a_{2m} \in \mathscr{A} \setminus \{a\}$,

$$\sum_{j=1}^{m} \frac{1}{a - a_j} = \sum_{j=m+1}^{2m} \frac{1}{a - a_j}$$

only if $(a_1,\dots,a_m)$ *and* $(a_{m+1},\dots,a_{2m})$ *are permutations of each other. Here, division (and addition) is taken in the field* $\mathbb{F}_p$.

(b) For every $\varepsilon>0$*, there exists* $P=P(\varepsilon)$ *such that for every* $p\geq P$*, every subset* $S\subseteq\mathscr{B}(p)$ *with* $|S|\geq p^{\ell}$ *satisfies* $E(S,S)\leq p^{-\gamma+\varepsilon}|S|^3$.

Pick α *such that*

$$\Omega(p^{\alpha})\leq|\mathscr{A}(p)|\leq p^{\alpha},\qquad|\mathscr{B}(p)|\geq\Omega(p^{1-\alpha+\varepsilon'}) \tag{8}$$

for some $\varepsilon'>0$ *and every sufficiently large* p*. Pick* $\varepsilon_1>0$ *for which there exist* $\alpha_1,\alpha_2,\varepsilon,x,y>0$ *such that*

$$\varepsilon_1+\varepsilon<\alpha_1-\alpha-(4/3)x-\varepsilon, \tag{9}$$
$$\ell\leq 1/2+(4/3)x-\alpha_1+\varepsilon/2, \tag{10}$$
$$\varepsilon_1+\varepsilon<\gamma/4-y/4-\varepsilon, \tag{11}$$
$$\alpha_2\geq 9x+\varepsilon, \tag{12}$$
$$c_0y/8-(\alpha_1/4+9\alpha_2/8)/m\leq x/8-\alpha/4, \tag{13}$$
$$\varepsilon_1+\varepsilon<c_0y/8-(\alpha_1/4+9\alpha_2/8)/m, \tag{14}$$
$$my\leq\min\{1/2-\alpha_1,1/2-\alpha_2\}, \tag{15}$$
$$3\alpha_2-2\alpha_1\leq(2-c_0)my. \tag{16}$$

Here, $c_0=1/10430$ *is a constant from Proposition 2 in [7]. Then for sufficiently large* p*, the* $p\times|\mathscr{A}(p)||\mathscr{B}(p)|$ *matrix with columns*

$$u_{a,b}:=\frac{1}{\sqrt{p}}\Big(e^{2\pi i(ax^2+bx)/p}\Big)_{x\in\mathbb{F}_p}\qquad a\in\mathscr{A},b\in\mathscr{B}$$

satisfies $(p^{1/2+\varepsilon_1/2-\varepsilon''},\delta)$*-RIP for any* $\varepsilon''>0$ *and* $\delta<\sqrt{2}-1$*, thereby implying* ExRIP$[1/2+\varepsilon_1/2]$.

Let's briefly discuss the structure of the proof of this result. As indicated in Section 2, the method is to prove flat-RIP-type cancellations, namely that

$$S(A_1,A_2):=\sum_{\substack{a_1\in A_1\\ a_2\in A_2}}\sum_{\substack{b_1\in\Omega_1(a_1)\\ b_2\in\Omega_2(a_2)}}\left(\frac{a_1-a_2}{p}\right)e_p\left(\frac{(b_1-b_2)^2}{2(a_1-a_2)}\right) \tag{17}$$

has size $\leq p^{1-\varepsilon_1-\varepsilon}$ whenever Ω_1 and Ω_2 are disjoint with size $\leq\sqrt{p}$. (Actually, we get to assume that these subsets and the $\Omega_i(a_i)$'s satisfy certain size constraints since we have an extra $-\varepsilon$ in the power of p; having this will imply the general case without the ε, as made clear in the proof of Theorem 2.) This bound is proved by considering a few different cases. First, when the $\Omega_i(a_i)$'s are small, (17) is small by a triangle inequality. Next, when the $\Omega_i(a_i)$'s are large, then we can apply a triangle inequality over each A_i and appeal to hypothesis (b) in Theorem 2 and Lemma 7. However, this will only give sufficient cancellation when the A_i's are small. In the

remaining case, Bourgain et al. prove sufficient cancellation by invoking Lemma 10 in [5], which concerns the following quantity:

$$T_{a_1}(A_2,B) := \sum_{\substack{b_1\in B\\ a_2\in A_2,\ b_2\in\Omega_2(a_2)}} \left(\frac{a_1-a_2}{p}\right) e_p\left(\frac{(b_1-b_2)^2}{4(a_1-a_2)}\right). \tag{18}$$

Specifically, Lemma 10 in [5] gives that $|T_{a_1}(A_2,B)|$ is small whenever B has sufficient additive structure. In the proof of the main result, they take a maximal subset $B_0 \subseteq \Omega_1(a_1)$ such that $|T_{a_1}(A_2,B_0)|$ is small, and then they use this lemma to show that $\Omega_1(a_1)\setminus B_0$ necessarily has little additive structure. By Lemma 7, this in turn forces $|T_{a_1}(A_2,B_1)|$ to be small, and so $|T_{a_1}(A_2,\Omega_1(a_1))|$ (and furthermore $|S(A_1,A_2)|$) are also small due to a triangle inequality. The reader is encouraged to find more details in the proofs found in Section 5.

What follows is a generalized version of the statement of Lemma 10 in [5], which we then use in the hypothesis of a generalized version of Lemma 2 in [5]:

Definition 4. Let $\mathrm{L10} = \mathrm{L10}[\alpha_1,\alpha_2,k_0,k_1,k_2,m,y]$ denote the following statement about subsets $\mathscr{A} = \mathscr{A}(p)$ and $\mathscr{B} = \mathscr{B}(p)$ of $\mathbb{F}_p$ for p prime:

For every $\varepsilon > 0$, there exists $P > 0$ such that for every $p \geq P$ the following holds:

Take $\Omega_1,\Omega_2 \subseteq \mathscr{A}\times\mathscr{B}$ such that

$$|A_2| \geq p^y, \tag{19}$$

and for which there exist powers of two M_1,M_2 such that

$$\frac{M_i}{2} \leq |\Omega_i(a_i)| < M_i \tag{20}$$

and

$$|A_i|M_i \leq 2\sqrt{p} \tag{21}$$

for $i = 1,2$ and for every $a_i \in A_i$. Then for every $B \subseteq \mathbb{F}_p$ such that

$$p^{1/2-\alpha_1} \leq |B| \leq p^{1/2} \tag{22}$$

and

$$|B-B| \leq p^{\alpha_2}|B|, \tag{23}$$

we have that (18) satisfies

$$|T_{a_1}(A_2,B)| \leq |B|p^{1/2-\varepsilon_2+\varepsilon} \tag{24}$$

with $\varepsilon_2 = k_0 y - (k_1\alpha_1 + k_2\alpha_2)/m$ for every $a_1 \in A_1$.

Lemma 9 (generalized version of Lemma 2 in [5]). *Take $\mathscr{A}$ arbitrarily and $\mathscr{B}$ satisfying the hypothesis (b) in Theorem 2, pick α such that $|\mathscr{A}(p)| \leq p^\alpha$ for every sufficiently large p, and pick $\varepsilon,\varepsilon_1,x > 0$ such that* L10 *holds with* (9)–(12) *and*

$$\varepsilon_1 + \varepsilon < \varepsilon_2 \leq x/8 - \alpha/4. \tag{25}$$

Then the following holds for every sufficiently large p:

Take $\Omega_1, \Omega_2 \subseteq \mathscr{A} \times \mathscr{B}$ for which there exist powers of two M_1, M_2 such that (20) *and* (21) *hold for $i = 1, 2$ and every $a_i \in A_i$. Then* (17) *satisfies $|S(A_1, A_2)| \leq p^{1-\varepsilon_1-\varepsilon}$.*

The following result gives sufficient conditions for L10, and thus Lemma 9 above:

Lemma 10 (generalized version of Lemma 10 in [5]). *Suppose $\mathscr{A}$ satisfies hypothesis (a) in Theorem 2. Then* L10 *is true with $k_0 = c_0/8$, $k_1 = 1/4$ and $k_2 = 9/8$ provided* (15) *and* (16) *are satisfied.*

These lemmas are proved in Section 5. With these in hand, we are ready to prove the main result:

Proof (Proof of Theorem 2). By Lemma 10, we have that L10 is true with

$$\varepsilon_2 = c_0 y/8 - (\alpha_1/4 + 9\alpha_2/8)/m.$$

As such, (13) and (14) together imply (25), and so the conclusion of Lemma 9 holds. We will use this conclusion to show that the matrix identified in Theorem 2 satisfies $(p^{1/2}, p^{-\varepsilon_1})$-weak flat RIP. Indeed, this will imply $(p^{1/2}, p^{-\varepsilon_1/2})$-flat RIP by Lemma 2, $(p^{1/2}, 75p^{-\varepsilon_1/2}\log p)$-RIP by Lemma 1, and $(p^{1/2+\varepsilon_1/2-\varepsilon''}, 75p^{-\varepsilon''}\log p)$-RIP for any $\varepsilon'' > 0$ by Lemma 3 (taking $s = p^{\varepsilon_1/2-\varepsilon''}$). Since $75p^{-\varepsilon''}\log p < \sqrt{2}-1$ for sufficiently large p, this will prove the result.

To demonstrate $(p^{1/2}, p^{-\varepsilon_1})$-weak flat RIP, pick disjoint $\Omega_1, \Omega_2 \subseteq \mathscr{A} \times \mathscr{B}$ of size $\leq p^{1/2}$. We need to show

$$\left| \left\langle \sum_{(a_1,b_1)\in\Omega_1} u_{a_1,b_1}, \sum_{(a_2,b_2)\in\Omega_2} u_{a_2,b_2} \right\rangle \right| \leq p^{1/2-\varepsilon_1}.$$

Recall that

$$\begin{aligned}
&\sum_{(a_1,b_1)\in\Omega_1} \sum_{(a_2,b_2)\in\Omega_2} \langle u_{a_1,b_1}, u_{a_2,b_2}\rangle \\
&\quad = \sum_{\substack{a_1\in A_1 \\ a_2\in A_2}} \sum_{\substack{b_1\in\Omega_1(a_1) \\ b_2\in\Omega_2(a_2)}} \langle u_{a_1,b_1}, u_{a_2,b_2}\rangle \\
&\quad = \sum_{\substack{a_1\in A_1 \\ a_2\in A_2\setminus A_1}} \sum_{\substack{b_1\in\Omega_1(a_1) \\ b_2\in\Omega_2(a_2)}} \frac{\sigma_p}{\sqrt{p}} \left(\frac{a_1-a_2}{p}\right) e_p\left(-\frac{(b_1-b_2)^2}{4(a_1-a_2)}\right).
\end{aligned}$$

As such, we may assume that A_1 and A_2 are disjoint without loss of generality, and it suffices to show that

$$\left| \sum_{\substack{a_1\in A_1 \\ a_2\in A_2}} \sum_{\substack{b_1\in\Omega_1(a_1) \\ b_2\in\Omega_2(a_2)}} \left(\frac{a_1-a_2}{p}\right) e_p\left(-\frac{(b_1-b_2)^2}{4(a_1-a_2)}\right) \right| \leq p^{1-\varepsilon_1}.$$

For each k, define the set

$$A_i^{(k)} := \{a_i \in A_i : 2^{k-1} \le |\Omega_i(a_i)| < 2^k\}.$$

Then we have

$$|A_i^{(k)}|2^{k-1} \le \sum_{a_i \in A_i^{(k)}} |\Omega_i(a_i)| = |\{(a_i,b_i) \in \Omega_i : a_i \in A_i^{(k)}\}| \le |\Omega_i| \le p^{1/2}.$$

As such, taking $M_i = 2^k$ gives that $A_i \leftarrow A_i^{(k)}$ satisfies (20) and (21), which enables us to apply the conclusion of Lemma 9. Indeed, the triangle inequality and Lemma 9 together give

$$\begin{aligned}\bigg|\sum_{\substack{a_1\in A_1\\ a_2\in A_2}}\sum_{\substack{b_1\in\Omega_1(a_1)\\ b_2\in\Omega_2(a_2)}}\bigg(\frac{a_1-a_2}{p}\bigg)e_p\bigg(-\frac{(b_1-b_2)^2}{4(a_1-a_2)}\bigg)\bigg| &\le \sum_{k_1=1}^{\lceil \frac{1}{2}\log_2 p\rceil}\sum_{k_2=1}^{\lceil \frac{1}{2}\log_2 p\rceil}|S(A_1^{(k_1)},A_2^{(k_2)})|\\ &\le p^{1-\varepsilon_1-\varepsilon}\log^2 p\\ &\le p^{1-\varepsilon_1}\end{aligned}$$

for sufficiently large p. $\square$

To summarize Theorem 2, we may conclude ExRIP$[1/2+\varepsilon_1/2]$ if we can find

(i) $m \in 2\mathbb{N}$ satisfying hypothesis (a),
(ii) $\ell, \gamma > 0$ satisfying hypothesis (b),
(iii) α satisfying (8), and
(iv) $\alpha_1, \alpha_2, \varepsilon, x, y > 0$ satisfying (9)–(16).

Since we want to conclude ExRIP$[z]$ for the largest possible z, we are inclined to maximize ε_1 subject to (i)–(iv), above. To find m, ℓ, γ, α which satisfy (i)–(iii), we must leverage a particular construction of $\mathscr{A}$ and $\mathscr{B}$, and so we turn to Lemma 8 and Corollary 2. Indeed, for any given m, Lemma 8 constructs $\mathscr{A}$ satisfying hypothesis (a) such that

$$\alpha = 1/(2m(4m-1)) \tag{26}$$

satisfies the first part of (8). Next, if we take $\beta := \alpha - \varepsilon'$ and define $r := \lfloor \beta \log_2 p \rfloor$ and $M := 2^{1/\beta - 1}$, then (2) constructs $\mathscr{B}$ which, by Corollary 2, satisfies hypothesis (b) provided

$$(2\tau - 1)(\ell - \gamma) \ge 10\gamma, \tag{27}$$

where τ is the solution to

$$\Big(\frac{1}{M}\Big)^{2\tau} + \Big(\frac{M-1}{M}\Big)^{\tau} = 1.$$

For this construction, $|\mathscr{B}| = M^r \ge \Omega(p^{1-\beta})$, thereby satisfying the second part of (8).

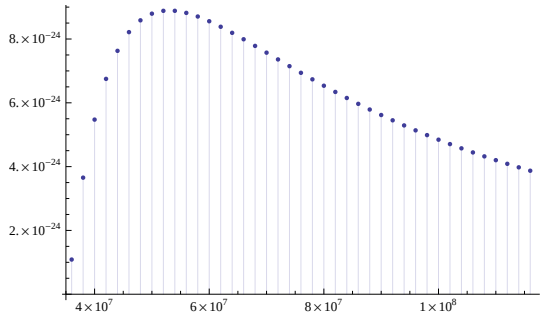

Fig. 1 The supremum of ε_1 as a function of m. Taking $\varepsilon' = 0$, we run a linear program to maximize ε_1 subject to the closure of the constraints (9)–(16), (27) for various values of m. A locally maximal supremum of $\varepsilon_1 \approx 8.8933 \times 10^{-24}$ appears around $m = 53,000,000$.

It remains to maximize ε_1 for which there exist $m, \varepsilon', \ell, \gamma, \alpha_1, \alpha_2, \varepsilon, x, y$ satisfying (9)–(16), (26) and (27). Note that m and ε' determine α and τ, and the remaining constraints (9)–(16), (27) which define the feasible tuples $(\varepsilon_1, \ell, \gamma, \alpha_1, \alpha_2, \varepsilon, x, y)$ are linear inequalities. As such, taking the closure of this feasibility region and running a linear program will produce the supremum of ε_1 subject to the remaining constraints. This supremum increases monotonically as $\varepsilon' \to 0$, and so we only need to consider the limiting case where $\varepsilon' = 0$. Running the linear program for various values of m reveals what appears to be a largest supremum of $\varepsilon_1 \approx 8.8933 \times 10^{-24}$ at $m = 53,000,000$ (see Fig. 1). Dividing by 2 then gives a new record:

$$\varepsilon_0 \approx 4.4466 \times 10^{-24}.$$

While this optimization makes a substantial improvement (this is over 8,000 times larger than the original record of Bourgain et al. in [5]), the constant is still tiny! For this particular construction of $\mathscr{A}$ and $\mathscr{B}$, the remaining bottlenecks may lie at the very foundations of additive combinatorics. For example, if $c_0 = 1/2$, then taking $m = 10,000$ leads to ε_0 being on the order of 10^{-12}.

5 Proofs of Technical Lemmas

This section contains the proofs of the technical lemmas (Lemmas 9 and 10) which were used to prove the main result (Theorem 2).

5.1 Proof of Lemma 9

First note that $|A_1|M_1 < p^{1/2+(4/3)x+\alpha-\alpha_1+\varepsilon}$ implies that

$$|A_1||\Omega_1(a_1)| < p^{1/2+(4/3)x+\alpha-\alpha_1+\varepsilon}$$

by (20), and by (21), we also have

$$|A_2||\Omega_2(a_2)| < 2p^{1/2}.$$

As such, the triangle inequality gives that

$$|S(A_1,A_2)| \le |A_1||A_2||\Omega_1(a_1)||\Omega_2(a_2)| \le 2p^{1+(4/3)x+\alpha-\alpha_1+\varepsilon} \le p^{1-\varepsilon_1-\varepsilon},$$

where the last step uses (9). Thus, we can assume $|A_1|M_1 \ge p^{1/2+(4/3)x+\alpha-\alpha_1+\varepsilon}$, and so the assumption $|\mathscr{A}| \le p^\alpha$ gives

$$M_1 \ge \frac{1}{|A_1|}p^{1/2+(4/3)x+\alpha-\alpha_1+\varepsilon} \ge p^{1/2+(4/3)x-\alpha_1+\varepsilon}. \tag{28}$$

Applying (20) and (28) then gives

$$|\Omega_1(a_1)| \ge \frac{M_1}{2} \ge \frac{1}{2}p^{1/2+(4/3)x-\alpha_1+\varepsilon} > p^{1/2+(4/3)x-\alpha_1+\varepsilon/2} \ge p^\ell,$$

where the last step uses (10). Note that we can redo all of the preceding analysis by interchanging indices 1 and 2. As such, we also have $|\Omega_2(a_2)| > p^\ell$. (This will enable us to use hypothesis (b) in Theorem 2.) At this point, we bound

$$\left|\sum_{\substack{b_1\in\Omega_1(a_1)\\ b_2\in\Omega_2(a_2)}} e_p\left(\frac{(b_1-b_2)^2}{4(a_1-a_2)}\right)\right|$$

using Lemma 7:

$$\le |\Omega_1(a_1)|^{1/2}E(\Omega_1(a_1),\Omega_1(a_1))^{1/8}|\Omega_2(a_2)|^{1/2}E(\Omega_2(a_2),\Omega_2(a_2))^{1/8}p^{1/8}.$$

Next, since $|\Omega_1(a_1)|, |\Omega_2(a_2)| > p^\ell$, hypothesis (b) in Theorem 2 with $\varepsilon \leftarrow 4\varepsilon$ gives

$$\le |\Omega_1(a_1)|^{7/8}|\Omega_2(a_2)|^{7/8}p^{1/8-\gamma/4+\varepsilon}.$$

At this point, the triangle inequality gives

$$\begin{aligned}|S(A_1,A_2)| &\le \sum_{\substack{a_1\in A_1\\ a_2\in A_2}}\left|\sum_{\substack{b_1\in\Omega_1(a_1)\\ b_2\in\Omega_2(a_2)}} e_p\left(\frac{(b_1-b_2)^2}{4(a_1-a_2)}\right)\right|\\ &\le \sum_{\substack{a_1\in A_1\\ a_2\in A_2}}|\Omega_1(a_1)|^{7/8}|\Omega_2(a_2)|^{7/8}p^{1/8-\gamma/4+\varepsilon}\end{aligned}$$

which can be further bounded using (20) and (21):

$$\le 2^{7/4}|A_1|^{1/8}|A_2|^{1/8}p^{1-\gamma/4+\varepsilon}$$

Thus, if $|A_1|, |A_2| < p^y$, then

$$|S(A_1,A_2)| \le 2^{7/4} p^{1+y/4-\gamma/4+\varepsilon} \le p^{1-\varepsilon_1-\varepsilon},$$

where the last step uses (11). As such, we may assume that either $|A_1|$ or $|A_2|$ is $\ge p^y$. Without loss of generality, we assume $|A_2| \ge p^y$. (Considering (19), this will enable us to use L10.)

At this point, take $B_0 \subseteq \Omega_1(a_1)$ to be a maximal subset satisfying (24) for $B \leftarrow B_0$, and denote $B_1 := \Omega_1(a_1) \setminus B_0$. Then the triangle inequality gives

$$|T_{a_1}(A_2,B_1)| \le \sum_{a_2 \in A_2} \left| \sum_{\substack{b_1 \in B_1 \\ b_2 \in \Omega_2(a_2)}} e_p\left(\frac{(b_1-b_2)^2}{4(a_1-a_2)}\right)\right|,$$

and then Lemma 7 gives

$$\le \sum_{a_2 \in A_2} |B_1|^{1/2} E(B_1,B_1)^{1/8} |\Omega_2(a_2)|^{1/2} E(\Omega_2(a_2),\Omega_2(a_2))^{1/8} p^{1/8}.$$

This can be bounded further by applying $E(\Omega_2(a_2),\Omega_2(a_2)) \le |\Omega_2(a_2)|^3$, (20), (21) and the assumption $|\mathscr{A}| \le p^\alpha$:

$$\le 2^{7/8} |B_1|^{1/2} E(B_1,B_1)^{1/8} p^{\alpha/8+9/16}. \tag{29}$$

At this point, we claim that $E(B_1,B_1) \le p^{-x} M_1^3$. To see this, suppose otherwise. Then $|B_1|^3 \ge E(B_1,B_1) > p^{-x} M_1^3$, implying

$$|B_1| > p^{-x/3} M_1, \tag{30}$$

and by (20), we also have

$$E(B_1,B_1) > p^{-x} M_1^3 > p^{-x} |\Omega_1(a_1)|^3 \ge p^{-x} |B_1|^3.$$

Thus, Lemma 5 with $K = p^x$ produces a subset $B_1' \subseteq B_1$ such that

$$|B_1'| \ge \frac{|B_1|}{20p^x} > \frac{M_1}{20p^{(4/3)x}} \ge \frac{1}{20} p^{1/2-\alpha_1+\varepsilon} \ge p^{1/2-\alpha_1}$$

where the second and third inequalities follow from (30) and (28), respectively, and

$$|B_1' - B_1'| \le 10^7 p^{9x} |B_1| \le p^{9x+\varepsilon} |B_1| \le p^{\alpha_2} |B_1|,$$

where the last step follows from (12). As such, $|B_1'|$ satisfies (22) and (23), implying that $B \leftarrow B_1'$ satisfies (24) by L10. By the triangle inequality, $B \leftarrow B_0 \cup B_1'$ also satisfies (24), contradicting B_0's maximality.

We conclude that $E(B_1,B_1) \le p^{-x} M_1^3$, and continuing (29) gives

$$|T_{a_1}(A_2,B_1)| \le 2^{7/8}|B_1|^{1/2}M_1^{3/8}p^{9/16+\alpha/8-x/8}.$$

Now we apply (24) to $B \leftarrow B_0$ and combine with this to get

$$\begin{aligned}|T_{a_1}(A_2,\Omega_1(a_1))| &\le |T_{a_1}(A_2,B_0)| + |T_{a_1}(A_2,B_1)| \\ &\le |B_0|p^{1/2-\varepsilon_1} + 2^{7/8}|B_1|^{1/2}M_1^{3/8}p^{9/16+\alpha/8-x/8}.\end{aligned}$$

Applying $|B_0|,|B_1| \le |\Omega_1(a_1)| \le M_1$ by (20) then gives

$$|T_{a_1}(A_2,\Omega_1(a_1))| \le M_1p^{1/2-\varepsilon_1} + 2^{7/8}M_1^{7/8}p^{9/16+\alpha/8-x/8}.$$

Now we apply the triangle inequality to get

$$\begin{aligned}|S(A_1,A_2)| &\le \sum_{a_1\in A_1}|T_{a_1}(A_2,\Omega_1(a_1))| \\ &\le |A_1|\Big(M_1p^{1/2-\varepsilon_1} + 2^{7/8}M_1^{7/8}p^{9/16+\alpha/8-x/8}\Big),\end{aligned}$$

and applying (21) and the assumption $|\mathscr{A}| \le p^\alpha$ then gives

$$\le 2p^{1-\varepsilon_2} + 2^{7/4}p^{1+\alpha/4-x/8} \le 2p^{1-\varepsilon_2} + 2^{7/4}p^{1-\varepsilon_2} \le p^{1-\varepsilon_1-\varepsilon},$$

where the last steps use (25). This completes the proof.

5.2 Proof of Lemma 10

We start by following the proof of Lemma 10 in [5]. First, Cauchy–Schwarz along with (20) and (21) give

$$\begin{aligned}|T_{a_1}(A_2,B)|^2 &= \bigg|\sum_{\substack{a_2\in A_2\\ b_2\in\Omega_2(a_2)}}\bigg(\frac{a_1-a_2}{p}\bigg)\cdot\sum_{b_1\in B}e_p\bigg(\frac{(b_1-b_2)^2}{4(a_1-a_2)}\bigg)\bigg|^2 \\ &\le 2\sqrt{p}\sum_{\substack{a_2\in A_2\\ b_2\in\Omega_2(a_2)}}\bigg|\sum_{b_1\in B}e_p\bigg(\frac{(b_1-b_2)^2}{4(a_1-a_2)}\bigg)\bigg|^2.\end{aligned}$$

Expanding $|w|^2 = w\overline{w}$ and applying the triangle inequality then gives

$$= 2\sqrt{p}\bigg|\sum_{\substack{a_2\in A_2\\ b_2\in\Omega_2(a_2)}}\sum_{b_1,b\in B}e_p\bigg(\frac{b_1^2-b^2}{4(a_1-a_2)} - \frac{b_2(b_1-b)}{2(a_1-a_2)}\bigg)\bigg| \le 2\sqrt{p}\sum_{b_1,b\in B}|F(b,b_1)|,$$

where

$$F(b,b_1) := \sum_{\substack{a_2\in A_2\\ b_2\in\Omega_2(a_2)}} e_p\left(\frac{b_1^2-b^2}{4(a_1-a_2)} - \frac{b_2(b_1-b)}{2(a_1-a_2)}\right).$$

Next, Hölder's inequality $\|F1\|_1 \le \|F\|_m\|1\|_{1-1/m}$ gives

$$|T_{a_1}(A_2,B)|^2 \le 2\sqrt{p}|B|^{2-2/m}\left(\sum_{b_1,b\in B}|F(b,b_1)|^m\right)^{1/m}. \tag{31}$$

To bound this, we use a change of variables $x := b_1+b \in B+B$ and $y := b_1-b \in B-B$ and sum over more terms:

$$\sum_{b_1,b\in B}|F(b,b_1)|^m \le \sum_{\substack{x\in B+B\\ y\in B-B}}\left|\sum_{\substack{a_2\in A_2\\ b_2\in\Omega_2(a_2)}} e_p\left(\frac{xy}{4(a_1-a_2)} - \frac{b_2y}{2(a_1-a_2)}\right)\right|^m.$$

Expanding $|w|^m = w^{m/2}\overline{w}^{m/2}$ and applying the triangle inequality then gives

$$= \left|\sum_{\substack{x\in B+B\\ y\in B-B}}\sum_{\substack{a_2^{(i)}\in A_2\\ b_2^{(i)}\in\Omega_2(a_2^{(i)})\\ 1\le i\le m}} e_p\left(\sum_{i=1}^{m/2}\left[\tfrac{xy}{4(a_1-a_2^{(i)})} - \tfrac{b_2y}{2(a_1-a_2^{(i)})} - \tfrac{xy}{4(a_1-a_2^{(i+m/2)})} + \tfrac{b_2y}{2(a_1-a_2^{(i+m/2)})}\right]\right)\right|$$

$$\le \sum_{y\in B-B}\sum_{\substack{a_2^{(i)}\in A_2\\ b_2^{(i)}\in\Omega_2(a_2^{(i)})\\ 1\le i\le m}}\left|\sum_{x\in B+B} e_p\left(\frac{xy}{4}\sum_{i=1}^{m/2}\left[\frac{1}{a_1-a_2^{(i)}} - \frac{1}{a_1-a_2^{(i+m/2)}}\right]\right)\right|.$$

Next, we apply (20) to bound the number of m-tuples of $b_2^{(i)}$'s for each m-tuple of $a_2^{(i)}$'s (there are less than M_2^m). Combining this with the bound above, we know there are complex numbers $\varepsilon_{y,\xi}$ of modulus ≤ 1 such that

$$\sum_{b_1,b\in B}|F(b,b_1)|^m \le M_2^m\sum_{y\in B-B}\sum_{\xi\in\mathbb{F}_p}\lambda(\xi)\varepsilon_{y,\xi}\sum_{x\in B+B}e_p(xy\xi/4), \tag{32}$$

where

$$\lambda(\xi) := \left|\left\{a^{(1)},\dots,a^{(m)}\in A_2 : \sum_{i=1}^{m/2}\left[\frac{1}{a_1-a^{(i)}} - \frac{1}{a_1-a^{(i+m/2)}}\right] = \xi\right\}\right|.$$

To bound the $\xi=0$ term in (32), pick $a^{(1)},\dots,a^{(m)}\in A_2$ such that

$$\sum_{i=1}^{m/2}\left[\frac{1}{a_1-a^{(i)}} - \frac{1}{a_1-a^{(i+m/2)}}\right] = 0. \tag{33}$$

Then

$$\sum_{i=1}^{m/2} \frac{1}{a_1 - a^{(i)}} + \frac{m}{2} \cdot \frac{1}{a_1 - a^{(1)}} = \sum_{i=m/2+1}^{m} \frac{1}{a_1 - a^{(i)}} + \frac{m}{2} \cdot \frac{1}{a_1 - a^{(1)}},$$

and so by hypothesis (a) in Theorem 2, we have that $(a^{(1)},\dots,a^{(m/2)},a^{(1)},\dots,a^{(1)})$ is a permutation of $(a^{(m/2+1)},\dots,a^{(m)},a^{(1)},\dots,a^{(1)})$, which in turn implies that $(a^{(1)},\dots,a^{(m/2)})$ and $(a^{(m/2+1)},\dots,a^{(m)})$ are permutations of each other. Thus, all possible solutions to (33) are determined by $(a^{(1)},\dots,a^{(m/2)})$. There are $|A_2|^{m/2}$ choices for this $m/2$-tuple, and for each choice, there are $(m/2)!$ available permutations for $(a^{(m/2+1)},\dots,a^{(m)})$. As such,

$$\lambda(0) = (m/2)!|A_2|^{m/2}, \tag{34}$$

which we will use later to bound the $\xi = 0$ term. In the meantime, we bound the remainder of (32). To this end, it is convenient to define the following functions:

$$\zeta'(z) := \sum_{\substack{y\in B-B\\ \xi\in\mathbb{F}_p^*\\ y\xi=z}} \varepsilon_{y,\xi}\lambda(\xi), \qquad \zeta(z) := \sum_{\substack{y\in B-B\\ \xi\in\mathbb{F}_p^*\\ y\xi=z}} \lambda(\xi).$$

Note that $|\zeta'(z)| \le \zeta(z)$ by the triangle inequality. We use the triangle inequality and Hölder's inequality to bound the $\xi \neq 0$ terms in (32):

$$\begin{aligned}
&\left|\sum_{y\in B-B}\sum_{\xi\in\mathbb{F}_p^*}\lambda(\xi)\varepsilon_{y,\xi}\sum_{x\in B+B}e_p(xy\xi/4)\right| \\
&\quad= \left|\sum_{\substack{x\in B+B\\ z\in\mathbb{F}_p}}\zeta'(z)e_p(xz/4)\right| \\
&\quad\le \sum_{x\in\mathbb{F}_p}\left|1_{B+B}(x)\cdot\sum_{z\in\mathbb{F}_p}\zeta'(z)e_p(xz/4)\right| \\
&\quad\le |B+B|^{3/4}\left(\sum_{x\in\mathbb{F}_p}\left|\sum_{z\in\mathbb{F}_p}\zeta'(z)e_p(xz/4)\right|^4\right)^{1/4}.
\end{aligned} \tag{35}$$

To proceed, note that

$$\begin{aligned}
\left(\sum_{z\in\mathbb{F}_p}\zeta'(z)e_p(xz/4)\right)^2 &= \sum_{z,z''\in\mathbb{F}_p}\zeta'(z)\zeta'(z'')e_p(x(z+z'')/4) \\
&= \sum_{z'\in\mathbb{F}_p}(\zeta'*\zeta')(z')e_p(xz'/4),
\end{aligned}$$

where the last step follows from a change of variables $z' = z + z''$. With this and Parseval's identity, we continue (35):

$$= |B+B|^{3/4}\bigg(\sum_{x\in\mathbb{F}_p}\bigg|\sum_{z'\in\mathbb{F}_p}(\zeta'*\zeta')(z')e_p(xz'/4)\bigg|^2\bigg)^{1/4}$$
$$= |B+B|^{3/4}\|\zeta'*\zeta'\|_2^{1/2}p^{1/4}$$
$$\leq |B+B|^{3/4}\|\zeta*\zeta\|_2^{1/2}p^{1/4}, \tag{36}$$

where the last step follows from the fact that $|(\zeta'*\zeta')(z)| \leq (\zeta*\zeta)(z)$, which can be verified using the triangle inequality. Since $\zeta(z) = \sum_{\xi\in\mathbb{F}_p^*} 1_{B-B}(z/\xi)\lambda(\xi)$, the triangle inequality gives

$$\begin{aligned}\|\zeta*\zeta\|_2 &= \bigg\|\bigg(\sum_{\xi\in\mathbb{F}_p^*}\lambda(\xi)1_{\xi(B-B)}\bigg)*\bigg(\sum_{\xi'\in\mathbb{F}_p^*}\lambda(\xi')1_{\xi'(B-B)}\bigg)\bigg\|_2\\ &\leq \sum_{\xi,\xi'\in\mathbb{F}_p^*}\lambda(\xi)\lambda(\xi')\|1_{\xi(B-B)}*1_{\xi'(B-B)}\|_2\\ &= \sum_{\xi,\xi'\in\mathbb{F}_p^*}\lambda(\xi)\lambda(\xi')\|1_{B-B}*1_{(\xi'/\xi)(B-B)}\|_2,\end{aligned} \tag{37}$$

where the last step follows from the (easily derived) fact that $1_{B-B}*1_{(\xi'/\xi)(B-B)}$ is a dilation of $1_{\xi(B-B)}*1_{\xi'(B-B)}$.

To bound (37), we will appeal to Corollary 2 in [5], which says that for any $A\subseteq\mathbb{F}_p$ and probability measure λ over $\mathbb{F}_p$,

$$\sum_{b\in\mathbb{F}_p^*}\lambda(b)\|1_A*1_{bA}\|_2 \ll (\|\lambda\|_2+|A|^{-1/2}+|A|^{1/2}p^{-1/2})^{c_0}|A|^{3/2}, \tag{38}$$

where $\ll$ is Vinogradov notation; $f\ll g$ means $f=O(g)$. As such, we need to construct a probability measure and understand its 2-norm. To this end, define

$$\lambda_1(\xi) := \frac{\lambda(\xi)}{\|\lambda\|_1} = \frac{\lambda(\xi)}{|A_2|^m}. \tag{39}$$

The sum $\sum_{\xi\in\mathbb{F}_p}\lambda(\xi)^2$ is precisely the number of solutions to

$$\frac{1}{a_1-a^{(1)}}+\cdots+\frac{1}{a_1-a^{(m)}}-\frac{1}{a_1-a^{(m+1)}}-\cdots-\frac{1}{a_1-a^{(2m)}}=0,$$

which by hypothesis (a) in Theorem 2, only has trivial solutions. As such, we have

$$\|\lambda\|_2^2 = m!|A_2|^m. \tag{40}$$

At this point, define $\lambda_1'(b)$ to be $\lambda_1(\xi'/b)$ whenever $b\neq 0$ and $\lambda_1(0)$ otherwise. Then λ_1' is a probability measure with the same 2-norm as λ_1, but it allows us to directly apply (38):

$$\sum_{\xi\in\mathbb{F}_p^*}\lambda_1(\xi)\|1_{B-B}*1_{(\xi'/\xi)(B-B)}\|_2$$
$$=\sum_{b\in\mathbb{F}_p^*}\lambda_1'(b)\|1_{B-B}*1_{b(B-B)}\|_2$$
$$\ll(\|\lambda_1\|_2+|B-B|^{-1/2}+|B-B|^{1/2}p^{-1/2})^{c_0}|B-B|^{3/2}. \quad (41)$$

At this point, our proof deviates from the proof of Lemma 10 in [5]. By (39), (40) and (19), we have

$$\|\lambda_1\|_2=|A_2|^{-m}\|\lambda\|_2\le\sqrt{m!}|A_2|^{-m/2}\le\sqrt{m!}p^{-my/2},$$

Next, (22) and (23) together give

$$|B-B|\ge|B|\ge p^{1/2-\alpha_1}$$

and

$$|B-B|\le p^{\alpha_2}|B|\le p^{1/2+\alpha_2}.$$

Thus,

$$\|\lambda_1\|_2+|B-B|^{-1/2}+|B-B|^{1/2}p^{-1/2}\le\sqrt{m!}p^{-my/2}+p^{\alpha_1/2-1/4}+p^{\alpha_2/2-1/4}$$
$$\le p^{-my/2+4m\varepsilon}, \quad (42)$$

where the last step follows from (15). So, by (37), (39), (41) and (42), we have

$$\|\zeta*\zeta\|_2\le|A_2|^{2m}\sum_{\xi'\in\mathbb{F}_p^*}\lambda_1(\xi')\sum_{\xi\in\mathbb{F}_p^*}\lambda_1(\xi)\|1_{B-B}*1_{(\xi'/\xi)(B-B)}\|_2$$
$$\ll|A_2|^{2m}(\|\lambda_1\|_2+|B-B|^{-1/2}+|B-B|^{1/2}p^{-1/2})^{c_0}|B-B|^{3/2}$$
$$\le|A_2|^{2m}p^{-(c_0/2)my+4c_0m\varepsilon}|B-B|^{3/2}, \quad (43)$$

and subsequent application of (32), (34), (36) and (43) gives

$$\sum_{b_1,b\in B}|F(b,b_1)|^m\le(\tfrac{m}{2})!(M_2|A_2|)^m|A_2|^{-m/2}|B-B||B+B|$$
$$+O(M_2^m|A_2|^m|B-B|^{3/4}|B+B|^{3/4}p^{-(c_0/4)my+2c_0m\varepsilon}p^{1/4}). \quad (44)$$

By Lemma 4 in [5] (which states that $|A+A|\le|A-A|^2/|A|$), condition (23) implies

$$|B+B|\le\frac{|B-B|^2}{|B|}\le p^{2\alpha_2}|B|.$$

We now use this with (21), (19) and (23) to bound (44):

$$\ll(\tfrac{m}{2})!(2\sqrt{p})^mp^{-my/2}p^{3\alpha_2}|B|^2+(2\sqrt{p})^mp^{(9/4)\alpha_2}|B|^{3/2}p^{-(c_0/4)my+2c_0m\varepsilon}p^{1/4}$$

Next, the left-hand inequality of (22) gives that $p^{1/4} \leq |B|^{1/2} p^{\alpha_1/2}$, leading to the following bound:

$$\ll |B|^2 p^{m/2-my/2+3\alpha_2} + |B|^2 p^{m/2+\alpha_1/2+(9/4)\alpha_2-(c_0/4)my+2c_0m\varepsilon}.$$

Overall, we have

$$\sum_{b_1,b\in B} |F(b,b_1)|^m \leq 2^{-m}|B|^2 p^{m/2+\alpha_1/2+(9/4)\alpha_2-(c_0/4)my+2m\varepsilon},$$

since $c_0 < 1$, and $3\alpha_2 - 2\alpha_1 \leq (2-c_0)my$ (i.e., (16)). Thus, (31) gives

$$\begin{aligned} |T(A_2,B)|^2 &\leq \sqrt{p}|B|^{2-2/m}(|B|^2 p^{m/2+\alpha_1/2+(9/4)\alpha_2-(c_0/4)my+2m\varepsilon})^{1/m} \\ &= |B|^2 p^{1-(c_0y/4-\alpha_1/(2m)-(9\alpha_2)/(4m))+2\varepsilon}. \end{aligned}$$

Finally, taking square roots produces the result.

Acknowledgements The author was supported by NSF Grant No. DMS-1321779. The views expressed in this chapter are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

References

1. L. Applebaum, S. D. Howard, S. Searle, R. Calderbank, Chirp sensing codes: Deterministic compressed sensing measurements for fast recovery, Appl. Comput. Harmon. Anal. 26 (2009) 283–290.
2. A. S. Bandeira, E. Dobriban, D. G. Mixon, W. F. Sawin, Certifying the restricted isometry property is hard, IEEE Trans. Inform. Theory 59 (2013) 3448–3450.
3. A. S. Bandeira, M. Fickus, D. G. Mixon, P. Wong, The road to deterministic matrices with the restricted isometry property, J. Fourier Anal. Appl. 19 (2013) 1123–1149.
4. R. Baraniuk, M. Davenport, R. DeVore, M. Wakin, A simple proof of the restricted isometry property for random matrices, Constr. Approx. 28 (2008) 253–263.
5. J. Bourgain, S. J. Dilworth, K. Ford, S. Konyagin, D. Kutzarova, Explicit constructions of RIP matrices and related problems, Duke Math. J. 159 (2011) 145–185.
6. J. Bourgain, M. Z. Garaev, On a variant of sum-product estimates and explicit exponential sum bounds in prime fields, Math. Proc. Cambridge Philos. Soc. 146 (2009) 1–21.
7. J. Bourgain, A. Glibichuk, Exponential sum estimate over subgroup in an arbitrary finite field, Available online: `http://www.math.ias.edu/files/avi/Bourgain_Glibichuk.pdf`
8. E. J. Candès, The restricted isometry property and its implications for compressed sensing, C. R. Acad. Sci. Paris, Ser. I 346 (2008) 589–592.
9. P. G. Casazza, M. Fickus, Fourier transforms of finite chirps, EURASIP J. Appl. Signal Process. 2006 (2006).
10. R. A. DeVore, Deterministic constructions of compressed sensing matrices, J. Complexity 23 (2007) 918–925.
11. M. Fickus, D. G. Mixon, J. C. Tremain, Steiner equiangular tight frames, Linear Algebra Appl. 436 (2012) 1014–1027.
12. S. Foucart, H. Rauhut, A Mathematical Introduction to Compressive Sensing, Berlin, Springer, 2013.

13. P. Koiran, A. Zouzias, Hidden cliques and the certification of the restricted isometry property, Available online: arXiv:1211.0665
14. D. G. Mixon, Deterministic RIP matrices: Breaking the square-root bottleneck, Short, Fat Matrices (weblog) http://dustingmixon.wordpress.com/2013/12/02/deterministic-rip-matrices-breaking-the-square-root-bottleneck/
15. D. G. Mixon, Deterministic RIP matrices: Breaking the square-root bottleneck, II, Short, Fat Matrices (weblog) http://dustingmixon.wordpress.com/2013/12/11/deterministic-rip-matrices-breaking-the-square-root-bottleneck-ii/
16. D. G. Mixon, Deterministic RIP matrices: Breaking the square-root bottleneck, III, Short, Fat Matrices (weblog) http://dustingmixon.wordpress.com/2014/01/14/deterministic-rip-matrices-breaking-the-square-root-bottleneck-iii/
17. L. R. Welch, Lower bounds on the maximum cross correlation of signals, IEEE Trans. Inform. Theory 20 (1974) 397–399.
18. T. Tao, Open question: deterministic UUP matrices, What's new (weblog) http://terrytao.wordpress.com/2007/07/02/open-question-deterministic-uup-matrices/
19. T. Tao, V. H. Vu, Additive Combinatorics, Cambridge U. Press, 2006.